

WHITE PAPER

System Virtualization: Sun Microsystems Enables Choice, Flexibility, and Management

Sponsored by: Sun Microsystems

John Humphreys
October 2006

EXECUTIVE SUMMARY

Virtualization is quickly becoming an important technology across all parts of the IT environment. The technology is rapidly being incorporated in storage, networks, and client environments. By far the most visible adoption of virtualization technology is happening in servers — from the largest Unix server to the smallest volume system. Virtualization software breaks the link between a given software-driven application or service and the hardware resources of the underlying system and encapsulates and isolates that stack or service. It is being used as a tool for consolidation and more recently as a means to bring high availability, disaster recovery, and resource balancing to a larger part of organizations' IT infrastructures.

Sun is emphasizing choice in how companies virtualize. The company is developing, partnering, and making technology available so that virtualization can evolve rapidly and deliver more value to organizations. By focusing on innovation, both through its own development and by making technology available to the community, Sun is working to ensure that users have access to a portfolio of virtualization tools to address a broad base of business requirements. Sun is also incorporating virtualization management into its N1 System Manager so that customers have a single pane of glass when managing their infrastructures. Sun's customer-driven approach for how companies organize and manage their virtual environments will enable customers to choose the right virtualization tool for the job yet still maintain cohesive management of a heterogeneous environment. This customer-centric approach will allow virtualization to become a standard deployment scenario for the enterprise.

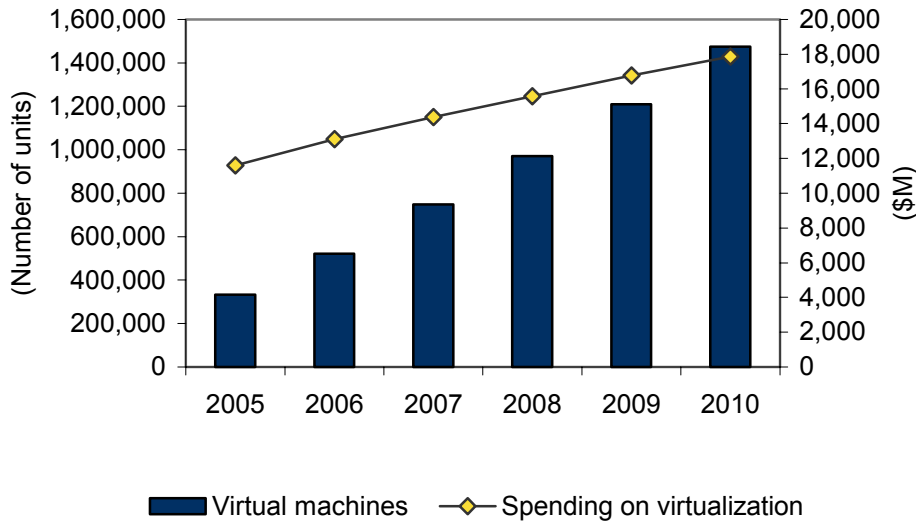
SITUATION OVERVIEW

It's been said that everything old is new again, and that saying is particularly appropriate for system virtualization. Virtualization and system partitioning have existed for decades aboard mainframe and Unix based systems. What's new is the concept of bringing this virtualization technology to small and medium-scale x86-based servers.

According to IDC, as shown in Figure 1, worldwide shipments of server hardware platforms with virtualization capabilities installed are currently forecast to grow rapidly from 330,000 units in 2005 to over 1,400,000 units in 2010, driving over \$18 billion in customer spending of hardware, software, and services in support of virtualized environments.

FIGURE 1

The Rise of Virtualization



Source: IDC, 2006

Although this data examines only the server market, it is important to view virtualization as a technology that spans the entire range of systems in an organization — not just servers but also storage, application, network, and client infrastructures.

Virtualize Everything!

Virtualization software breaks the link between — or "decouples" — a given software-driven application or service and the hardware resources of the underlying system. This decoupling means that functionality delivered by software may survive the unexpected loss of the original host systems. As these capabilities become more widely adopted, leading organizations are empowered to look at their IT infrastructures as a pool of shared resources to drive utilization, agility, and resiliency across all IT services. The following descriptions highlight each of the core segments of the virtualization market:

Virtual access software. This software allows applications to be accessed from nearly any intelligent access point device over just about any network, without the applications having to have been architected to support that device or network.

To support a broad variety of computing needs, Sun is addressing this market with the Sun Secure Global Desktop suite as well as the Sun Ray thin client. The concept of thin clients has been around for a long time, but customers recently have begun to rethink their client virtualization strategies as the cost of desktop support as well as data and device losses have raised security concerns within organizations. These concerns, coupled with new performance and price innovations, suddenly make virtual clients a cost-justifiable reality for a much broader swath of the enterprise.

Virtual application environment. This software creates an application development and deployment environment that allows properly developed applications to be more robust and reliable and also unaware of the underlying operating environments and hardware platforms. These benefits are available only to applications written for this environment. Application server software and parallel database software fit in this category.

Because Sun developed Java and because Java is a staple of application servers, the company is extremely active in the application server market. Not only is Sun the biggest contributor to the code base, but it has made much of the software freely available and has articulated plans to open-source the Java stack to incent its broadest possible use. Application server environments are one of the largest virtualization categories today with acceptance in all market segments aboard all major operating systems (OSs) and hardware platforms.

Virtual processing software. This category of software ranges from virtual machine (VM) software making a single system appear to be many systems, each supporting its own operating environment, to single-system image clustering software, which makes many systems appear to be a single computing resource running a single operating environment. This category also includes parallel processing software, load-balancing software, and data and application availability software.

Server virtualization, a segment of virtual processing, is one of the most diverse and rapidly growing parts of the virtual environment today. The relevant technologies are very broad and range from hardware partitions and traditional resource management to virtual machines and virtualizing within an operating system. Sun offers Dynamic System Domains, Solaris Containers, and Solaris Resource Manager as well as partnerships with the major virtual machine vendors to ensure customers a best-of-breed solution stack. Highlighting the differences in the technologies and the benefits and challenges associated with each option is the thrust of this paper.

Virtual storage. This software allows applications to be unaware of where and how application and data files are actually stored. It virtualizes not only within a particular vendor's storage arrays but also across multiple vendors' arrays so that users can set up a single pool of storage no matter their hardware installed base.

In the storage virtualization space, Sun is offering StorageTek Virtual Storage Manager (VSM) system, which is a virtual tape solution. The company also has the StorageTek 6920 and 9990, which allow IT organizations to consolidate multiple disk systems into a single 9990 system.

Server provisioning and management. This software makes it possible for operators and administrators to load, manage, and operate multisystem configurations regardless of whether any of the other virtual environment software categories are present. It is a primary component of on-demand or adaptive environment approaches to application deployment.

In terms of provisioning and management, Sun's N1 System Manager nests with the N1 Service Provisioning System to enable users to enable to monitor, manage, and provision systems, applications, and service levels across the datacenter. Details on how Sun is extending this product to manage virtual environments are covered later in the document.

Server Virtualization: The Need for Choice

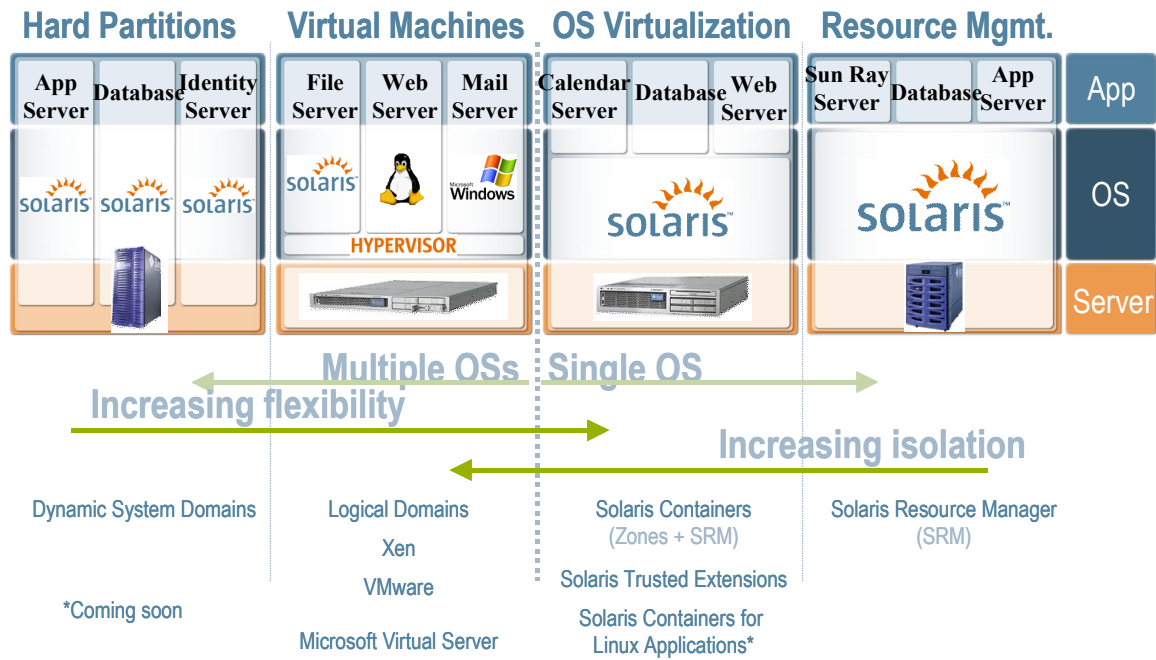
In essence, virtualization is a hardware resource-sharing strategy. The core advantage of implementing server virtualization is the ability to drive higher utilizations on physical servers by sharing the use of hardware resources across several workloads, moving away from the "one application per server" approach that often characterized the use of nonvirtualized servers in the past. Another advantage is the ability to isolate one workload from another, in terms of memory, data, and hard drive contents, avoiding such problems as driver conflicts when two applications are running on the same server, and to provide better security and minimize compliance issues while sharing resources across multiple workloads.

As the operational and business benefits delivered by virtualization have been realized, implementing virtualization has become a major initiative for organizations. There is growing recognition among customers of the need for diversity in how they virtualize or decouple the application stack from the underlying hardware. This need for diversity stems directly from the application as different applications have different requirements in terms of isolation, flexibility, and performance. Having the right virtualization tool for the job is as critical as it is for choosing a hardware platform or vendor.

Sun has responded by offering its customers choices at the hardware, operating system, and hypervisor layers. In this manner, customers can continue to choose not only the application that best meets their needs but also the virtualization software, operating system, and hardware that will support those needs (see Figure 2). It should also be noted that while these technologies are presented independently, they can be combined to enhance the capabilities and counteract some of the challenges of each standalone option.

FIGURE 2

Spectrum of Virtualization Solutions from Sun



Source: Sun Microsystems, 2006

Types of Server Virtualization

As mentioned earlier in the document, several different types of virtualization technologies are relevant in the market today, including:

Hard partitions. This well-established technology has been around in mainframe and RISC Unix systems for quite some time. Sun has had hard partitions for over a decade. The workload isolation goes down to the hardware level and is based on the specific number of processor boards — whereby one or more boards can be dedicated to an application and a separate image of the operating system. Use of hard partitions is limited by the number of processor boards in a system; the technology is static and not very granular.

From a hardware partitioning perspective, Sun continues to offer customers Dynamic System Domains in its larger SPARC-based systems. This technology has long been used in support of running multiple mission-critical applications on the same host, yet with complete software isolation. (IBM offers LPARs and HP has nPars.) The technology enables the user to take down (for maintenance or upgrades) one partition while ensuring the other partitions on the machine continue to operate. It also allows the customer to dynamically resize a partition while the system is running. This approach to partitioning is a staple in the RISC Unix world and continues to advance in terms of the granularity of the partitions and the degree of flexibility with which the customer can manage them.

Resource management. This well-established and mature technology offers some of the same value propositions as virtualization. Resource management allows for the allocation of system resources across multiple applications. This level of product is available from multiple vendors. IBM has Workload Manager and HP offers Global Workload Manager, while Sun has long had Solaris Resource Manager for scheduling and allocating resources within a single OS instance. Unlike hard partitions, resource management involves the installation of only a single image of the operating system with multiple applications installed on top of the OS. A resource scheduler controls access to all the system resources. The technology is very flexible and granular, but the applications are not truly isolated from each other; therefore, contention issues can come into play and all the applications must ultimately run on the same OS. Sun makes this type of virtualization solution available on both its RISC and x86 systems.

Operating system virtualization. This virtualization option isolates applications in a single OS. Sun has Solaris Containers and SWsoft has Virtuozzo. Solaris Containers, which are part of Solaris, combine the isolation benefits of Solaris Zones with the flexibility benefits of Solaris Resource Manager. Containers are freely available to customers with an upgrade to Solaris 10. With OS virtualization, the virtualization is instrumented in the OS and the applications reside in isolated containers in a single copy of the operating system. This type of virtualization is very lightweight and requires almost no system or application overhead. Any application faults are isolated in the container and do not impact the other zones (though if the OS or server fails, so do all the containers, illustrating the importance of having a robust operating environment and hardware). OS virtualization also scales with the operating system and can run on the entire line of Sun servers (both x86 and SPARC). Like resource management, each container can access all the resources on the system so that applications have burst capacity to handle spikes in demand. This approach also helps consolidate OS images so that changes and upgrades to core operating environments are greatly streamlined.

With container technology, users can partition Solaris into two or more environments. One environment contains a "global zone" for administration, patching, scripting and resource control, and the other zones are available to run the user software. This single OS approach helps customers address concerns around virtual machine or OS sprawl. It also allows administrators to manage more applications than with the "one server, one application" approach so typical today. In the future, one can almost envision a container running across the entire datacenter.

Because containers run in Solaris 10, customers also can leverage the security of Solaris Trusted Extensions to ensure applications and data are isolated even when they reside on the same OS. With Solaris 10, users are able to label containers and control access at the kernel level to ensure a secure and solid access method even when public and private data is on the same host.

Sun is working on a BrandZ as an extension to Solaris Containers technology. BrandZ, which will be productized as Solaris Containers for Linux Applications, enables containers to run different Solaris and Linux applications unmodified. This approach will allow customers to consolidate multiple operating environments into a container while still running a single OS. It also will allow Linux users to take advantage of Solaris tools such as DTrace.

Additionally, Sun is investing in container technology so that customers receive more fine-grained resource management, including capping of CPU resources and controlling memory. The company is working to offer greater flexibility in container security so that privileges can be adjusted per container. Sun will also add container mobility in the future. This capability, which has proved so compelling for VMware users, will allow a container to be moved from one physical host to another. It will allow containers to quickly attach and detach from hosts, allow containers to be cloned, and offer a streamlined means of backing up and restoring. These advancements are expected in November 2006.

Virtual machines. A virtual machine is a collection of different solutions that virtualize below the operating system, whereby a thin layer of software enables a customer to run multiple identical or different operating systems and applications in isolation from each other. Like containers, a "virtual machine" can isolate faults so that the surrounding VMs can continue to run. Unlike containers, each VM requires a separate operating system so that while hardware consolidation is possible, there is no reduction in the number of OS images. Additionally, some overhead is associated with running the virtualization software. This overhead is negligible for many applications, but for I/O-intensive applications, it can become a significant concern. Today there are two different types of virtual machines: software virtualized and paravirtualized.

☒ **Software virtualization** must intercept some OS calls and service those using an emulation of underlying hardware resources, leading to higher overhead. Software virtualization traps privileged instructions that are used by an operating system to control hardware functions and handles them through emulation. This approach is used by VMware and by HP in Integrity Virtual Machines.

☒ **Paravirtualization** is similar to pure software virtualization, but it offers lower overhead to the guest operating system. The key difference between software virtualization and paravirtualization is that a paravirtualized environment mandates that the privileged instructions be replaced or trapped on a static basis, such that a specialized version of the operating system kernel must be used. This solution typically requires a kernel modification to the operating system before it can run. IDC notes that altering kernels may require some level of compliance testing to verify post-modification operation. This approach is used by Xen and the open source community. Sun is using a Xen hypervisor combined with a Solaris control domain to virtualize Solaris, Linux, and Windows software stacks. Sun feels the combination of Solaris with a Xen hypervisor provides a more stable and scalable virtualization solution for customers.

☐ Sun is leveraging paravirtualization in its Logical Domains. This approach can be compared with micro-partitions, such as those offered by IBM. Logical Domains have a lightweight hypervisor in firmware for Sun UltraSPARC T1 and future CMT systems. This hypervisor will allow users to create isolated, full virtual machine domains to consolidate multiple applications on a single server. Initially, Sun will enable Solaris, Linux, and FreeBSD software stacks to run in Logical Domains. Users will be able to manage the virtual machines independently, including being able to start, stop, and reboot a VM. In Solaris 10, users will be able to dynamically add and remove virtual CPUs while the OS is still running to resize the VM on the fly.

From a pure performance perspective, paravirtualization would offer a performance edge compared with pure software virtualization, while from configuration and compatibility perspectives, pure software virtualization is easier to adopt. However, absolute performance issues aside, the specific type of virtual machine solution being used on a server will be invisible to users who are interacting with application software.

In terms of software-based or hypervisor partitioning, Sun x64 systems work with all the major implementations, most notably those from VMware (ESX Server and VMware Server [formerly GSX Server]), Microsoft (Microsoft Virtual Server), and Xen.

Manageability. Sun has made integrated manageability a central theme across its entire product line. The company has taken a nested approach that starts with common on-board service processors (ILOM) and extends to the infrastructure with N1 System Manager. At the application and service levels, N1 System Manager nests with N1 Service Provisioning System to cover management from the most granular component to the broadest datacenter service.

Customers can use the N1 System Manager for infrastructure management across the entire Sun server portfolio. This software package allows users to manage groups of servers, provision operating system patches and firmware updates remotely, monitor hardware and operating system events, as well as have lights-out and role-based access to the systems. N1 System Manager supports and works with third-party solutions from certified operating system vendors including Solaris, Red Hat, and Novell SUSE Linux and with virtualization products from VMware and Xen.

Data from N1 System Manager also feeds into the N1 Service Provisioning System so that administrators can focus on managing not just infrastructure but, more importantly, the services IT provides to end users. Increasingly, service management is becoming the focus of large organizations; therefore, automation in the management of services and the underlying infrastructure is a key feature. The N1 Service Provisioning System is the vehicle for Sun to enable its customers to begin to deliver automation at all levels in the technology solution stack.

Why Virtualize Systems? The Business Case

At the highest level, IDC has found two major benefits associated with the virtualization of system resources: cost benefits and greater flexibility in managing systems. Because of the benefits, in some cases, virtualized systems are becoming the standard deployment platform for all new applications in customer environments.

In terms of cost benefits, we find virtualization impacts both capital and operational costs. Partitioning allows users to increase the server utilization of the current systems in their environments. Data suggests that system utilization averages between 5% and 25% in most situations, which means that 75% to 95% of the capacity is unused. This information, combined with the fact that many x86 customers still assign one server per application, underscores the reality that many customers have dramatically overprovisioned their hardware resources. Virtualization is a means to increase utilization by consolidating multiple applications on a single host so that capital costs are used efficiently. This is one of the major reasons that most customers initially started to virtualize their servers.

From an operational perspective, virtualization is impacting costs in a few ways. By separating the application from the underlying hardware, IT administrators are able to more easily manage, provision, restart, and migrate applications on a shared pool of server hardware. IT is finding that virtualization is a means to speed change in the datacenter. Companies using virtualization technologies report that the time to deliver a new server within their organizations has been reduced from days or weeks to just a few hours or, in some cases, only minutes. This ability provides IT with a means to be more responsive to business requirements.

Customers also report massive increases in application-to-administrator ratios, which significantly drive down the cost of managing large IT infrastructures. This ability to grow infrastructures without adding staff members can lead to competitive advantages for businesses.

In addition, the reduction in physical hosts associated with better utilization of resources can lead to dramatic savings in electricity and cooling costs. Not only is this savings recurring, but the reduction in power consumption can prolong datacenter life, thus helping organizations avoid having to expand or build new facilities.

Server Virtualization Use Cases

Because of the benefits of virtualization, including the capital and operational cost benefits and the flexibility and ease of managing applications, users are developing an increasingly broad set of usage scenarios for virtualization technologies in their environments. These major scenarios are as follows:

- ☒ **Server consolidation.** The ability to run multiple production applications while sharing physical hardware resources leads to reduced hardware costs (and depending on the solution, potentially lower software costs), easier administration, and simplified manageability. Not only does virtualization help to reduce server footprints, but it also can greatly reduce power consumption, cooling demands, and cabling and free up space in overfilled datacenters.
- ☒ **Test and development.** A popular use is to run pilot solutions in different virtual partitions. This approach takes advantage of all the benefits of server consolidation (reducing the number of "nonproductive" servers) and allows increased agility through accelerated application rollout. New applications can be validated inside partitions before being rolled out across the datacenter. Additionally, developers can easily compile and test their applications under multiple operating systems using a single server, without heavy hardware investments.
- ☒ **Hosting legacy applications.** Virtualization allows users to run applications along with older operating systems that are no longer supported and that can no longer run on today's servers. The business case for legacy applications is that customers benefit from the increased performance of newer hardware without the cost of having to port the application off unsupported legacy operating environments.

- ☒ **Hardware migrations and upgrades.** Virtualization allows a staged migration of applications onto new hardware. By validating the solutions in virtual partitions, IT shops reduce the possibility of disruption of service due to migration. When migrating to a new architecture, staging the migration in virtualized servers allows robust validation before a widespread rollout. The ROI with migration comes from the level of automation and reduced IT staff time needed to move applications across platforms.
- ☒ **Business continuity.** Most IT shops today deploy some form of failover, usually involving replicated servers. Because many failures are associated with software, users are recognizing that replicating servers in virtual partitions on the same machine provides the benefits of higher application availability from built-in virtualization failover at a reduced hardware cost.
- ☒ **Capacity planning.** Virtual partitions can be sized and resized as required. This capability offers great flexibility to IT shops, which can configure large compute resources when required (e.g., end of the quarter), then scale down as needed. By scaling compute resources as needed, IT shops can optimize the overall utilization of the server. In the future, the provisioning of compute resources will become even more dynamic, providing the ability to add and remove resources as required.
- ☒ **Resource management.** Resource management refers to policy-based tools that monitor utilization and move virtual machines as needed to balance peak capacity with headroom. The benefit of load balancing is that users can more effectively pool their infrastructures and share a centralized managed pool across multiple applications. This ability to balance loads and provision applications dynamically helps further reduce hardware costs and increase service levels and application availability.

This ever-expanding portfolio of usage for virtualization is what makes the technology so compelling to end users and, in turn, is why Sun, its partners, and leading virtualization vendors are working together to create a robust platform that encompasses the hardware, software, and management layers of the solution. This collaborative approach is critical if the full potential of the technology is to be realized. When choosing a virtualization solution, customers must examine not just the cost but also their business needs, the application, the operating environment, the level of isolation, and the flexibility and performance they require. Only through a thorough examination of these criteria can users strike the right balance and find the best virtualization solution for their specific needs.

Getting Started with Virtualization

The adoption of virtualization technology within organizations typically follows a cycle of low-risk to mission-critical solutions, producing a corresponding low-value to high-value return to organizations. The typical adoption curve moves through a continuum of uses that includes test and development, production server consolidation, virtualization to support of business continuity, resource management, and eventually a utility computing environment.

Today, most organizations adopting virtualization technology on a broad scale are still in the middle of consolidating the excessive server sprawl that took place over the past decade. Consolidating several servers onto a single box is one of the key drivers of the first wave of virtualization.

Because consolidation in test, development, and production environments is the primary focus of organizations' virtualization activities today, Sun has developed a "virtual" consolidation service. This engagement is designed to use existing architecture and operations assessment as a basis for defining a new consolidation architecture for an organization's systems, network, and storage — leveraging the portfolio of virtualization technologies in Sun's portfolio. The engagement is composed of a consolidation workshop, business review, as well as consolidation design and implementation services.

CHALLENGES/OPPORTUNITIES

Sun and other vendors must continue to engage customers in industrywide educational discussions that articulate how to map IT requirements to the different virtualization, operating environment, and server platform choices available. Having a broad set of choices is the first step in full coverage of customer requirements.

In addition to choice, vendors must be able to offer the tools to manage both virtual and physical infrastructures. Tool proliferation is one of the many issues that inhibit streamlining IT. Sun is working to deliver more complete virtualization management within its N1 System Manager and N1 Service Provisioning System products. This work will need to continue to develop for virtualization to become even more pervasive.

Sun will need to stay ahead of the curve as virtualization moves beyond simply being a tool for consolidation. IDC expects the company to develop assessment, implementation, and support expertise to help customers take advantage of the new emerging use case for virtualization technologies. Vendors will need to differentiate these services as customers begin to look to transform and incorporate virtualization across their IT environments.

CONCLUSION

Virtualization is rapidly becoming a standard piece of the deployment platform in IT organizations because the technology provides tangible benefits both in terms of capital cost reductions and operational benefits. Additionally, the technology has already found a broad set of usages, ranging from improved hardware utilization in test and development, to application life-cycle extension, to high availability and disaster recovery.

The decoupling of the application from the underlying infrastructure also enables a greater degree of agility. The combination of a variety of use cases, demonstrable benefits, and increased agility is driving broad-scale adoption in the market.

But as with any broadly accepted technology, users require choice to make sure they have the right tool for the job. With virtualization, users must consider the degree of isolation, the breadth of operating systems they plan to support, the goals of the organization (Is server consolidation the only goal, or is a reduction in OS instances important?), the degree of mobility and flexibility the solution necessitates, the applications they plan to support, and the level of performance they require.

Sun's investments in Solaris Containers, Dynamic System Domains, Logical Domains, Solaris Resource Manager, and Xen, as well as its partnerships with key software virtualization vendors, illustrate how the company is making choice in virtualization a top priority. Sun is also integrating virtualization management into N1 System Manager and N1 Service Provisioning System so that it can provide customers a single pane of glass from which users will be able to manage physical and virtual systems from the infrastructure through the services provided on those systems. The company has also worked to create assessment and implementation programs that help customers process the litany of choices and emerge with a highly successful virtualization implementation.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2006 IDC. Reproduction without written permission is completely forbidden.